RATIONALE AND STRATEGY FOR VR STANDARDS

Dr Barry J Shepherd
IBM Corporation
Advanced Workstation and Systems Division
11400 Burnet Road
Austin, TX USA

The virtual reality industry needs to start developing a set of formal standards for several aspects of virtual reality (VR) systems and applications. This paper explains why standards are needed, the types of standards which are needed, and how they could be developed. It will draw on my experience as chairman of the ISO/IEC committee responsible for standards for computer graphics and the presentation of multimedia information.

Who wants standards, and why do they want them? In the absence of even defacto standards, the existence of many competing products can confuse customers who wish to acquire VR hardware or software. Thus, customers like to see standards because they simplify the buying decision. Some government agencies even limit purchases to those products which conform to relevant standards. Standards also promote interworking of components obtained from different vendors, and permit substituting one vendor's product for that of another. Within this paper, "standards" mean formal standards developed by ISO, or by an ANSI accredited group, since this assures they were developed by consensus, with open participation of affected parties, and formal responses to public review comments.

Standards also benefit developers, since they stabilize an area of technology for at least a couple of years. In general (and certainly in the area of computer graphics) standards define interfaces, not internal structure. Thus, a "smarter" vendor can provide a "better" (quicker or smaller) implementation and still conform to the standard. A robust standard is also extensible, so that (optional) enhancements can be provided by vendors who wish to satisfy special customer needs.

A set of standards to support various aspects of virtual reality systems will stabilize the market, and encourage greater participation by system and application providers, and by customers interested in acquiring a multivendor solution. The following discussion of areas for potential standards work is based on today's understanding of virtual reality. It can be expected to evolve over the next three years, which is the probable lead time before standards development starts. Also, it should be noted that several of the proposed areas of standardization are not strictly within the domain of traditional computer graphics. However, all these areas need to be standardized by some group in order to obtain a full system solution.

0-7803-1363-1/93 $3.00 © 1993 IEEE

AREAS FOR VR STANDARDS

Work is needed in two areas to support unencumbered input. One are a is voice recognition commands. That is, what are the default phrases to navigate in a virtual world, or effect some action on its contents? Different virtual worlds would probably tend to acquire their own specific vocabulary. However, some general principals would be very helpful. The second area is a gestural language, for moving in or modifying a virtual world. Does a participant go up by pointing up (which finger?) or by moving their whole hand upwards? Once again, only a set of defaults would be standardized, but general guidelines would be of great help. The gestures could be captured by a glove, or by a TV camera. With suitable real time pattern recognition, hand gestures thus could provide unencumbered input.

People who enter a virtual world require not only a way to indicate where they want to go, but also the equivalent of a map and compass to show them where they are. The industry needs to agree on common navigational aids for use within a virtual world. While it is important for a participant to know where they are in a virtual world, knowing how they got there is vital if they wish to leave in an orderly fashion. It may be inelegant or even somewhat embarrassing to have to use an escape/panic mechanism to leave a virtual world. However, that escape sequence should also be standardized, to make virtual reality less threatening. Becoming "lost" while studying a hypermedia document can be frustrating, but when a participant's surroundings are not only foreign (or perhaps even alien) but also seem real the need for good standardized navigational tools should be apparent.

Sound is an important part of a virtual world. However, to be effective, the sound must be modified to take the position and head orientation of the virtual world participant into account. The relative intensity of a virtual sound source must be represented by simultaneous presentation through several physical loudspeakers. Also, the sound from each speaker must be filtered to account for the frequency dependence of the directional sensitivity of the ear. If this processing is done by hardware, the interface to that hardware must be defined. An eventual standard must address specification of earprints (frequency dependent spatial sensitivity functions) for the participant, and how to describe the location of not only the virtual but also the physical sound sources (loud-speakers).

An extensible interface definition mechanism is needed to handle th e unusual input devices often required by virtual reality systems. It must support gloves, wands, head and body position trackers, and exo-skeleton joint angle values. A corresponding interface specification is needed for output devices such as tactile pads or force feedback devices. Not included in this initial list are heaters, coolers, wind machines, vibrators, odor synthesizers, or other more exotic devices.

The highly interactive nature of virtual reality has prompted the creation of special purpose operating systems. Some of these systems handle multiple cooperating processes and processors. Within ISO there is only one operating system being standardized, and that is Posix in SC22 WG15. This work, originally, started within the IEEE Computer Society, is addressing real time extensions, but it is not clear if it will

meet the needs of virtual reality applications. A standard set of interfaces for interprocess communication and resource allocation/ management would be helpful, but it is not obvious which standards committee would be able to work on this project.

When there are standards for the default ways to act in, and interact with, virtual reality and the required hardware interfaces have been standardized there is one additional area where standards can be of help. The virtual realities which are to be explored must be created. Please note that no one proposes that the content of those realities be standardized. However, in the same way that authoring systems can aid with the administrative issues involved with creating multimedia "documents," toolkits to aid the construction of virtual worlds are needed.

The toolkit specifications would most likely define a rich set of capabilities, which would be subset and adapted to the needs of particular application areas. This raises the question of the format used to store the definition of a virtual world, or reusable parts of it, which clearly need to be standardized. Not only the geometry of the world but also both the static and dynamic properties (such as density, rigidity, center of gravity, temperature, specific heat, etc) of its components must be stored. If an object oriented approach is used, then work of the OMG (Object Management Group) should be considered.

ISO STANDARDS PROCESS

If it is agreed that standards may be useful "in the future," then let us consider how they are developed. Information Technology standards are developed by Joint Technical Committee 1 (JTC1) which reports to both ISO (the International Standards Organization) and the IEC (the International Electrotechnical Committee). There are several subcommittees (SCs) in JTC1. JTC1 must authorize any new standards development activity undertaken by one of its SCs by means of a New work item Proposal (NP) letter ballot sent to the countries which participate on JTC1. They are asked to comment on the suitability of the proposal, and their willingness to help to staff the activity. The proposal must carefully state the goals and benefits of the activity, other groups requiring liaison, and a development schedule. For an activity in an evolving area of new technology, a study period prior to NP submission is allowed.

A standard evolves through a series of stages, starting as an Initial Draft (ID). After one or more iterations within the development group as a Working Draft (WD) it is approved by the responsible subcommittee for circulation as a Committee Draft (CD). A CD is supposed to be technically complete, but may still generate many technical comments. Unless the amount of technical change necessitated by the comments forces another CD ballot, the revised document is circulated for comment as a Draft International Standard (DIS). Any technical changes resulting from comments on the DIS will cause another DIS circulation, but normally the comments on a DIS are editorial so that the final International Standard (IS) text can be created for publication.

There is an alternative path used to process an existing specification as a standard. The specification can be submitted to ISO by some national body for "fast

track" processing. It is circulated for ballot as a DIS. It is circulated "as is," but if approved in the ballot, it must be reformatted to meet ISO rules. The processing of any comments received is assigned to an appropriate subcommittee. If there are too many NO votes, which could not be resolved without substantial technical change to the document, the document must be restarted as a new project. The fast track procedure does not allow a 2nd DIS ballot.

There are two common sources for fast track documents; consortia and defacto standards. Examples of consortia related to multimedia are the Interactive Multimedia Association (IMA) for multimedia interchange format and services standards, and the Musical Instrument Digital Interface (MIDI) for interfacing instruments and synthesizer/generators through computers. Some consortia charge membership fees sufficiently high to cover the development of a trial or sample implementation which is made available to the consortium members. This is not done by SCs.

Examples of defacto standards include PostScript from Adobe, the Tagged Image File Format (TIFF) from Aldus, and Windows from Microsoft. Defacto standards can bring stability to some area of technology, but cannot generally be viewed as open standards. However, there is a great range in the behavior of the owners of defacto standards. Adobe was an active contributor to the development of SPDL (the Standardized Page Description Language) in SC18 WG8. This standard is similar, but not identical, to Adobe's PostScript.

SC24 STANDARDS PROCESS

SC24, the committee which I chair, employs the JTC1 development procedures, as well as some additional practices of its own. All SC24 interchange and application programming interface (API) standards separate the definition of semantics from syntax. The functionality is not only device and system independent, it is also application independent. A "language binding" specifies the functionality of a "semantic standard" in the syntax, style, and data types of a particular programming language. One or more language bindings (to C, Fortran, Ada, etc) or encodings must be available at most one stage "behind" the semantic standard (ie, WD binding for CD semantic specification, CD for DIS, and DIS for final IS text). Working initially from a semantic standard is much easier than trying to recover the functional description from a standard which combines syntax and semantics.

User requirements are gathered before starting to define a standard . These requirements are used to measure the extent to which a standard addresses a problem defined by its community of users. A well defined set of user requirements can also stop "creeping featurism" which can slow the progress of any technical activity. Of course, it is possible to revise the user requirements in a controlled manner during the development process. To gather user requirements, a questionnaire was circulated to people who could be potential users of the standard. The other source of requirements was open meetings in conjunction with conferences. These meetings also provided the opportunity to seek new participants on the committee. Both techniques involve biased sample sets, since they accessed people with sufficient interest in the subject to con-

tribute their thoughts. However, these people are most likely to be early adopters of the standard, so that is the right kind of sample to get.

Even before drafting is started on a standard, there are technical choices to be made. SC24 uses "issue processing" to help with this task. When someone asks, "Should the standard do/contain X?" and the answer is not quickly and unanimously agreed to by the group, an issue is added to the issue log. A fully document issue contains a statement of the question, a list of alternative answers, and a series of arguments for and against the alternatives and arguments already listed. The technical developers then vote on the issue, seeking consensus on one of the alternatives. The issue log can be used to educate new participants and to shorten re-discussion of previously agreed matters. With a large set of agreed issues, the direction and approach of the first draft is essentially defined. Issues are added as required during the writing of the text, and even to help resolve some comments.

Because some standards can provide a rich array of functionality, they may not appeal to users with more modest or focussed requirements. The current solution to this problem is the use of Internationally Standardized Profiles (ISPs). These are agreed subsets of a standard, and are often developed by interested communities of users. To ensure that the subsets (profiles) are "complete," SC24 is starting to include rules for defining profiles within its standards. This solves half of the application oriented customization problem. The other half is addressed by permitting registration of added values for parameters/attributes or even new types of operations. Any use of these extension mechanisms reduces portability, but retains the education and terminology benefits from using a common "core."

Defining a standard is helpful, but most people use implementations of standards, not the standards themselves. Thus, is must be possible to test if an implementation truly reflects the capabilities specified by the standard. All SC24 standards must be described in a way that ensures they are testable. The actual test suites are provided by groups outside of SC24. Within the USA, NIST (the National Institute of Standards and Technology) administers validation testing and manages certification of graphics implementations for government use. As an aid to testing, SC24 API standards must include sample programs. Language bindings of semantic standards must include the type definitions and constants in compilable order to assist in syntax checking and to guarantee the validity of the syntax (if not the semantics) of the interface. Each language binding must also include bindings of at least the sample programs from the semantic standard.

LOCATION FOR VR STANDARDS DEVELOPMENT

It is instructive to note the connection between multimedia and VR. To be truly engaging, virtual reality requires not only computer graphics but also audio and, for some applications, video. However, in these areas VR is simply "faster, better" multimedia. Just like multimedia, VR applications can be classified as communications based or document based. The VR equivalent of a document based multimedia application is one person exploring a virtual world. Communications based VR applications

fall into two broad classes. One is telepresence systems, where several people share a common virtual meeting space. The purpose could be discussion, education (via simulation) or competition. This is an extension of existing Computer Supported Cooperative Work applications. It is also, at least philosophically, an extension of desk top video conferencing systems. The other major class of communications based VR is teleoperator systems.

Given the relationship between multimedia and VR, where could standards be developed for some aspects of virtual reality? I believe that JTC1 SC24 is the logical place to develop several of the standards needed to support VR. Virtual Reality is visually oriented, and SC24 has developed APIs for graphics systems. We are now working on a new Presentation Environment for Multimedia Objects (PREMO). This API project is object oriented, device independent, user extensible, and includes audio, video, and time as data types. It is structured to permit installation of different modelling systems (such as physically based models) and viewing systems (such as stereo image generators).

The Computer Graphics Metafile (CGM) is widely used to store and interchange geometry information, and is being extended to include user specifiable attributes, such as structure delimiters. A new metafile project would certainly consider requirements for storing and exchanging VR models, including their additional attributes. The current work on a toolkit to support image processing applications and interchange (IPI) will provide valuable experience in creating a toolkit to support VR world building applications. Thus, SC24 can provide many of the basic standards needed by the software side of the VR industry. However, SC24 needs your user requirements, so that in two or three years we can develop standards project proposals that will be so attractive you decide to join our committee and work on them.

CONCLUSIONS

I believe an integrated family of standards must be developed, for the sake of the VR industry and its customers. This will bring stability to the marketplace, thus increasing customer acceptance and encouraging developers to provide additional applications. Everyone wins from this approach. During its history, SC24 has learned how to produce quality standards for computer graphics. The techniques it acquired are being used to improve the standards it is currently developing. Because of the visual orientation of virtual reality, and its conceptual relationship to multimedia presentation, SC24 is a natural location for development of at least some of the standards related to VR. There are problems to be overcome, for example in getting adequate staffing to produce standards in a timely manner. Perhaps consortia can provide part of an innovative solution to this problem, and allow SC24 to meet the needs of the VR community.